REPORT DOCUMENTATION PAGE	Form Approved OMB NO. 0704-0188

The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggesstions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA, 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any oenalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.
PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.

1. REPORT DATE (DD-MM-YYYY)	2. REPORT TYPE	3. DATES COVERED (From - To)
09-07-2012	Final Report	13-Jul-2011 - 12-Apr-2012

4. TITLE AND SUBTITLE	5a. CONTRACT NUMBER
Detecting Deception in Multiscale Activity Graphs	W911NF-11-1-0273
	5b. GRANT NUMBER
	5c. PROGRAM ELEMENT NUMBER 1M30BM
6. AUTHORS Jon Doyle	5d. PROJECT NUMBER
	5e. TASK NUMBER
	5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAMES AND ADDRESSES	8. PERFORMING ORGANIZATION REPORT NUMBER
North Carolina State University Office of Contract and Grants Leazar Hall Lower Level- MC Raleigh, NC 27695 -7214	

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)	10. SPONSOR/MONITOR'S ACRONYM(S)
U.S. Army Research Office P.O. Box 12211 Research Triangle Park, NC 27709-2211	ARO
	11. SPONSOR/MONITOR'S REPORT NUMBER(S) 60322-NS-II.2

12. DISTRIBUTION AVAILABILITY STATEMENT

Approved for Public Release; Distribution Unlimited

13. SUPPLEMENTARY NOTES

The views, opinions and/or findings contained in this report are those of the author(s) and should not contrued as an official Department of the Army position, policy or decision, unless so designated by other documentation.

14. ABSTRACT

The problem studied in this research is that of detecting deception by participants in large-scale activities. The primary means pursued was that of obtaining a better understanding of how to describe and predict the flow of information through large populations of disparate types of participants. Such an understanding can be used to see if the observed results of deceptive information match what would be expected were the deception to originate from suspected sources. It can also be used to estimate susceptibility of participants or populations to deceptive

15. SUBJECT TERMS

deception, information flow, social networks

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	15. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			Jon Doyle
UU	UU	UU	UU		19b. TELEPHONE NUMBER 919-513-0423

Report Title

Detecting Deception in Multiscale Activity Graphs

ABSTRACT

The problem studied in this research is that of detecting deception by participants in large-scale activities. The primary means pursued was that of obtaining a better understanding of how to describe and predict the flow of information through large populations of disparate types of participants. Such an understanding can be used to see if the observed results of deceptive information match what would be expected were the deception to originate from suspected sources. It can also be used to estimate susceptibility of participants or populations to deceptive information originating from different sources.

Enter List of papers submitted or published that acknowledge ARO support from the start of the project to the date of this printing. List the papers, including journal references, in the following categories:

(a) Papers published in peer-reviewed journals (N/A for none)

<u>Received</u> <u>Paper</u>

TOTAL:

Number of Papers published in peer-reviewed journals:

(b) Papers published in non-peer-reviewed journals (N/A for none)

<u>Received</u> <u>Paper</u>

TOTAL:

Number of Papers published in non peer-reviewed journals:

(c) Presentations

Number of Presentations: 0.00

Non Peer-Reviewed Conference Proceeding publications (other than abstracts):

<u>Received</u> <u>Paper</u>

TOTAL:

Number of Non Peer-Reviewed Conference Proceeding publications (other than abstracts):

Peer-Reviewed Conference Proceeding publications (other than abstracts):

<u>Received</u> <u>Paper</u>

TOTAL:

Number of Peer-Reviewed Conference Proceeding publications (other than abstracts):

(d) Manuscripts

Received   Paper

TOTAL:

Number of Manuscripts:

Books

Received   Paper

TOTAL:

Patents Submitted

Patents Awarded

Awards

Graduate Students

NAME	PERCENT_SUPPORTED	Discipline
Andrew Wicker	0.20	
FTE Equivalent:	**0.20**	
Total Number:	**1**	

Names of Post Doctorates

NAME	PERCENT_SUPPORTED
FTE Equivalent:	
Total Number:	

Names of Faculty Supported

NAME	PERCENT_SUPPORTED	National Academy Member
Jon Doyle	0.28	No
FTE Equivalent:	**0.28**	
Total Number:	**1**	

Names of Under Graduate students supported

NAME	PERCENT_SUPPORTED
FTE Equivalent:	
Total Number:	

Student Metrics

This section only applies to graduating undergraduates supported by this agreement in this reporting period

The number of undergraduates funded by this agreement who graduated during this period:	0.00
The number of undergraduates funded by this agreement who graduated during this period with a degree in science, mathematics, engineering, or technology fields:......	0.00
The number of undergraduates funded by your agreement who graduated during this period and will continue to pursue a graduate or Ph.D. degree in science, mathematics, engineering, or technology fields:......	0.00
Number of graduating undergraduates who achieved a 3.5 GPA to 4.0 (4.0 max scale):......	0.00
Number of graduating undergraduates funded by a DoD funded Center of Excellence grant for Education, Research and Engineering:......	0.00
The number of undergraduates funded by your agreement who graduated during this period and intend to work for the Department of Defense	0.00
The number of undergraduates funded by your agreement who graduated during this period and will receive scholarships or fellowships for further studies in science, mathematics, engineering or technology fields:......	0.00

Names of Personnel receiving masters degrees

NAME	
Total Number:	

Names of personnel receiving PHDs

NAME	
Andrew Wicker	
Total Number:	**1**

Names of other research staff

NAME	PERCENT_SUPPORTED
FTE Equivalent:	
Total Number:	

Sub Contractors (DD882)

Inventions (DD882)

Scientific Progress

Technology Transfer

Final Report

Detecting Deception in Multiscale Activity Graphs

Army Research Office Reference W911NF-11-1-0273

Jon Doyle
North Carolina State University
Department of Computer Science
2298 EB-II, 890 Oval Drive, Campus Box 8206
Raleigh, NC 27695-8206

(919) 513-0423
doyle@csc.ncsu.edu

Contents

1 Statement of Problem Studied

The problem studied in this research is that of detecting deception by participants in large-scale activities. The primary means pursued was that of obtaining a better understanding of how to describe and predict the flow of information through large populations of disparate types of participants. Such an understanding can be used to see if the observed results of deceptive information match what would be expected were the deception to originate from suspected sources. It can also be used to estimate susceptibility of participants or populations to deceptive information originating from different sources.

2 Summary of the Most Important Results

Propagation of information, deceptive or not, has been studied in several contexts. In nonmathematical studies, social psychologists have studied mechanisms based on group conformity. Epidemiologists, computer scientists, electrical engineers and physicists have developed numerous mathematical models of different types of propagation processes, including voter and contact models, cascade models, random interacting networks, and diffusion models.

2.1 Modeling heterogeneous populations

These mathematical approaches treat populations as homogeneous in type and propagation mechanisms as unitary in form. In his doctoral dissertation [3], Andrew Wicker develops a richer formal model of information propagation and uses it to rephrase and extend earlier mathematical and sociological studies.

Wicker's model of information propagation permits one to identify and formalize different types of influence mechanisms, such as influence through group conformity or through authority. One formalizes actual influences in a population by identifying different relations among individuals and mechanism types operating over those relations. For example, different conformity mechanisms might act on the same individual through that individual's participation in different relations, such as his school friends or coworkers. Similarly, different authority mechanisms might operate through her professional society or employer.

Wicker's model of information propagation also permits one to identify and formalize types of individuals that differ in how they respond to information and influence. Each individual is subject to influence only through those relations in which they take part, so one individual can be subject to influences that another is not. Moreover, one individual can respond to multiple influences in a way that differs from the way another individual responds. Wicker formalizes these differences in behavioral types defined by the methods by which multiple influences are combined.

This approach thus allows formalization of populations containing individuals of heterogeneous types, involved in heterogeneous relations, and subject to heterogeneous mechanisms for transmission of information and influence.

2.2 Formal concepts

The formalization of heterogeneous populations and information transfer mechanisms begins with the notion of a *measurable population* $(\mathcal{X}, \Psi, \Im)$ characterized by a set of individuals $\mathcal{X}$, a mapping Ψ associating a set Ψ_x of individual mental states with each $x \in \mathcal{X}$, and a mapping $\Im$ associating a measurable space $(\Psi_x, \Im_x)$ with each $x \in \mathcal{X}$. This is just a simple assumption of a set of participants with possibly distinct finite or infinite state spaces. The only restrictive requirement is that one can form probability measures over each individual state space, and hence over the population as a whole.

An *influence mechanism* $m = (r_m, \hat{\Psi}_m, \Im_m, \phi_m, \mu_m)$ over a measurable population $(\mathcal{X}, \Psi, \Im)$ is characterized by

- A binary *mechanism relation* r_m over $\mathcal{X}$ that restricts possible influences to those that follow paths within the relation;
- A set $\hat{\Psi}_m$ of *mechanism states* that distinguish the beliefs or information possibly transmitted by the influence mechanism;
- An associated measure space of events $\Im_m$ over $\hat{\Psi}_m$ so that one can form probability measures over the set of mechanism states;
- A *mental state projection* mapping ϕ_m that associates an *individual mental state projection* function $\phi_m^x : \Psi_x \to \hat{\Psi}_m$ with each $x \in \mathcal{X}$ which interprets each possible mental state of x as determining a corresponding mechanism state; and

- A *mechanism measure* mapping μ_m that associates to each individual x an *individual measure assignment* function μ_m^x. Each mechanism relation r_m induces an *influence neighborhood* function $\delta_m : \mathcal{X} \rightarrow \mathit{Pwr}(\mathcal{X})$, and the measure assignment function μ_m^x identifies the transition probability measure $\mu_m^x[\hat{\psi}] : \Im_m \times \Im_m \rightarrow [0, 1]$ obtaining for each mechanism state $\hat{\psi} \in \phi_m(\mathit{\Psi}_{\delta_m(x)})$ of the influence neighborhood of x.

Note that this definition defines a mechanism in terms of a specific relation over a specific population. When the population and relation are taken as parameters, one obtains a *mechanism type*.

The transition probabilities given the mechanism state transition measure μ_m^X over $\phi_m(\mathit{\Psi}_X)$ induce a first-order Markov chain $(\hat{\psi}_m^t(X))_{t \geq 0}$ over population states.

Wicker's analysis treats the case in which individuals only change state under the influence of neighbors, that is, that isolated individuals do not change state spontaneously. The formalism itself requires no such restriction. Wicker also assumes that transition probabilities that do not change over time, meaning that the induced Markov chain is time-homogeneous. Wicker's analysis also focusses on the case of finite mechanism state spaces, which means that the induced Markov chains are finite-state.

This formalization is used to recast several types of influence mechanisms, including group conformity, voter models, authoritarian influences, contaigion, and preference change influenced by preference similarity.

As noted earlier, different individuals might respond differently to the same influences. The formalism characterizes such differences in behavior in terms of *mechanism combination methods*. A mechanism combination method over a measurable population $(\mathcal{X}, \mathit{\Psi}, \Im)$ is a function that takes a set M of influence mechanisms over the population to a *combined mechanism* $c(M) = (r_{c(M)}, \hat{\mathit{\Psi}}_{c(M)}, \Im_{c(M)}, \phi_{c(M)}, \mu_{c(M)})$ over $(\mathcal{X}, \mathit{\Psi}, \Im)$. The combined mechanism relation is just the union of the relations of the combined mechanisms, and the combined states, events, and probabilities are found as the product of the corresponding structures for the combined mechanisms.

This formalization suffices to express several simple combination mechanisms, including random combination, dictatorial combination, and convex combination (weighted averaging). Wicker's analysis focusses on convex combination methods.

Although the formalization permits formalization of heterogeneous populations in which different individuals employ different combination methods and are distributed throughout the population in different ways, Wicker's analysis focusses on homogeneous populations in which all individuals employ the same combination method.

2.3 Analytical results

This analytical framework allows formal statement of influence maximization and origin-localization problems.

In the influence maximization problem, one seeks to influence as large or valuable a subpopulation as possible through a target population for initial influence. There might be multiple possible targets that yield influence of the same value, in which case one seeks targets that minimize a cost measure. Wicker's analysis focusses on this problem, and employs target size as the cost measure. This means finding the smallest targets that produce the largest results.

In the origin-localization problem, one takes an observed result of presumed influence and seeks to identify possible origin populations that produce the observed results. Again, there might be multiple possible originas, in which case one seeks origins of minimal size that optionally maximize an probability distribution that describes expectations about the likelihood of different origins generating the results in question.

The first topic in addressing these questions is to focus attention on the cases in which influence mechanisms actually have effect, that is, on the subpopulations actually related by mechanism relations. In general there might be many isolated individuals not subject to any influence mechanism. In the case under study in which these do not change state, the only way to influence them is by including them in the targeted population. Although such individual targeting might be desirable, it is irrelevant to analyzing how influences propagate through inter-individual relations. The upshot is that the analysis begins by identifying topological properties of influence relations, namely the connected components of mechanism relations. The relation between relations of different mechanisms makes use of the notion of *bridges* between components of different relations. If bridges exist, then information propagating by one mechanism can cross over a bridge formed by another mechanism relation to influence more individuals than either mechanism separately.

Wicker's first results show that one can decompose and compose optimal targets in the case in which the multiple mechanisms have no bridges between them. That is, any optimal target for the whole population can be decomposed into its subsets within each component of the separate mechanism relations. More importantly, one can optimize within each component independently and take the union of these component targets to obtain a global optimal target. That is, one can look at each separate component, find an target within that component that maximizes influence within that component, and then take the union of all such targets to find an optimal target for the entire population.

The next set of results seeks to obtain additional information about optimal targets in case the influence mechanisms exhibit *population monotonicity*, that is, monotonically increasing or decreasing influence in the target population. Wicker shows that a mechanism is non-decreasing population monotonic if there is always a positive transition probability to each population state from a given population state, and that this condition can be checked locally by examining the transition tables for each individual. When one combines non-decreasing population monotonic mechanism, one obtains a combined mechanism in which the full population always maximizes influence, though it need not be optimal.

Wicker shows that authoritarian influence mechanisms, contagion mechanisms, conformity mechanisms, and voter model mechanisms are all non-decreasing population monotonic.

Another set of results considers the case in which influence mechanisms are *nonvanishing* in the sense that at least one individual in a targeted component retains the desired influence. A set of mechanisms is nonvanishing when each mechanism and their convex combination is nonvanishing. Wicker shows that one can bound below the size of optimal targets when the set of mechanisms is nonvanishing.

Wicker shows that contagion mechanisms are non-vanishing, but that authoritarian, conformity, and voter model mechanisms are not non-vanishing, that is, there exist cases in which the targeted influence disappears completely.

The interaction of mechanisms is analyzed in terms of formal notions of *mechanism interference*. One mechanism (or set of mechanisms) interferes positively with another when the influence obtained by adding the new mechanisms to the original meets or exceeds the influence obtained from the original set alone. Negative interference is defined similarly, and mechanism sets are *non-interfering* when they yield the same influence separately as together.

Wicker shows that contaigion mechanisms positively interefere with authoritarian mechanisms.

From these notions of interference, one defines the notion of *mechanism monotonicity*, in which the expected influence changes montononically as additional mechanisms are added. Wicker shows that sets of positively interfering mechanisms are non-decreasing mechanism monotone, and that sets of negatively interfering mechanisms are non-increasing mechanism monotone. He also shows that a set of non-interfering mechanisms are mechanism monotone.

These interference and monotonicity results are first steps toward a more comprehensive characterization of mechanism submodularity. Wicker describes some preliminary results in this direction at the end of his dissertation.

2.4 Other work

In addition to the work on information flow, portions of the effort studied methods for assessing the reliability of information reaching some participant by analyzing the sources and support of the information. The issue here is that deceptive information can persist among participants even after the originator of the information departs through a process of mutual reinforcement. Assessment of sources and support can be used to distinguish beliefs grounded in reliable information from self-supporting rumors.

The focus of this portion of the effort was on developing analysis methods based on the idea of mechanical wave propagation and the more general mechanical models of [1, 2]. In this approach, information connections between beliefs within and among individuals are modeled as nonmonotonic or probabilistic dependencies, and one regards changes in antecedent beliefs or mental attitudes as generating mechanical forces on consequent beliefs, which then generate further forces on the consequences of the consequences. These motions then superpose in the same way that independently generated ripples in a pond. The analysis of reliability seeks to obtain estimates of honesty by structural and probabilistic analysis of the equilibrium states or cyclic motions of the web of belief, including observations of the results of perturbations chosen to illuminate the sensitivity of the beliefs to ongoing deception.

Although some progress was made along these lines, work on the information flow analysis described in the preceding section eventually demanded most of the attention.

2.5 Publications

The first publication stemming from this work is the doctoral dissertation completed by Andrew Wicker [3]. This is available online through the NCSU library at the URL `http://www.lib.ncsu.edu/resolver/1840.16/7681`.

The second publication is in an article by Wicker and the PI [4] derived from this dissertation work. This is expected to be submitted very soon to the journal *Autonomous Agents and Multi-Agent Systems*.

References

[1] Jon Doyle. *Extending Mechanics to Minds: The Mechanical Foundations of Psychology and Economics*. Cambridge University Press, London, UK, 2006.

[2] Jon Doyle. Toward a quantitative theory of belief change: Structure, difficulty, and likelihood (a progress report). Technical Report TR-2010-20, Department of Computer Science, North Carolina State University, Raleigh, NC, September 2010.

[3] Andrew W. Wicker. *Leveraging Multiple Mechanisms for Information Propagation*. PhD thesis, North Carolina State University, Raleigh, NC, April 2012. http://www.lib.ncsu.edu/resolver/1840.16/7681.

[4] Andrew W. Wicker and Jon Doyle. Leveraging multiple mechanisms for information propagation. In preparation.